Importance sampling with transformed weights

Manuel A. Vázquez and Joaquín Míguez

The importance sampling (IS) method lies at the core of many Monte Carlo-based techniques. IS allows the approximation of a target probability distribution by drawing samples from a proposal (or *importance*) distribution, different from the target, and computing importance weights (IWs) that account for the discrepancy between these two distributions. The main drawback of IS schemes is the degeneracy of the IWs, which significantly reduces the efficiency of the method. It has been recently proposed to use transformed IWs (TIWs) to alleviate the degeneracy problem in the context of Population Monte Carlo, which is an iterative version of IS. However, the effectiveness of this technique for standard IS is yet to be investigated. In this letter we numerically assess the performance of IS when using TIWs, and show that the method can attain robustness to weight degeneracy thanks to a bias/variance trade-off.

arXiv:1702.01987v2 [stat.AP] 20 Apr 2017

Introduction: One classical application of Monte Carlo (MC) methods is the approximation of a distribution of interest (often referred to as *target* distribution) by means of random samples. In many practical situations it is not possible or convenient to draw samples directly from the target. In such a case, it is common to rely on the importance sampling (IS) principle [1]. It consists in drawing samples from a *proposal* distribution, which are then assigned importance weights (IWs) to compensate for the mismatch between the target and the proposal.

A critical drawback of the IS methodology is the degeneracy of the IWs. This happens when only a few samples have non-negligible IWs. Since samples with weights close to zero are irrelevant when building a Monte Carlo approximation, weight degeneracy reduces the efficiency of the method. This problem aggravates when performing inference in high-dimensional systems [2].

The term Population Monte Carlo (PMC) [3] refers to a class of iterative IS algorithms in which samples drawn from a proposal distribution are used to obtain a refined proposal that can be sampled again. In the context of PMC, the authors of [4] propose to apply a non-linear transformation to the IWs that reduces their variability, and hence alleviates the degeneracy problem. Although a theoretical analysis of the asymptotic convergence of the method is provided in [4], to date there is no published numerical assessment of the performance of the importance samplers with transformed IWs (TIWs) in a non-iterative setting and with finite sample size. Moreover, the impact of certain parameters that are relevant to the performance of the method has not been investigated. In this work we tackle these open issues.

Standard Monte Carlo: Let $\boldsymbol{\theta}$ be an unknown $K \times 1$ random vector with known prior density $p(\boldsymbol{\theta})$. Our goal is the Bayesian estimation of $\boldsymbol{\theta}$ given an $N \times 1$ vector of observations, $\mathbf{y}$, that relates to the former through a likelihood function, $p(\mathbf{y}|\boldsymbol{\theta})$. Specifically, we aim at approximating the posterior probability density function (pdf) of $\boldsymbol{\theta}$, i.e.,

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}) \tag{1}$$

using a collection of M random samples, $\Theta^M = \{\boldsymbol{\theta}^{(i)}\}_{i=1}^M$, in the space of $\boldsymbol{\theta}$. From the latter, it is easy to approximate any expectation of the form $\mathbb{E}_{p(\boldsymbol{\theta}|\mathbf{y})}\left[f(\boldsymbol{\theta})\right] = \int f(\boldsymbol{\theta})p(\boldsymbol{\theta}|\mathbf{y})d\boldsymbol{\theta}$, where $f : \mathbb{R}^K \to \mathbb{R}$ is some real integrable function of $\boldsymbol{\theta}$. For instance, the posterior mean of $\boldsymbol{\theta}$ can be approximated as $\mathbb{E}_{p(\boldsymbol{\theta}|\mathbf{y})}\left[f(\boldsymbol{\theta})\right] \approx \frac{1}{M}\sum_{i=1}^M \boldsymbol{\theta}^{(i)}$.

Importance sampling: Let us denote by $\pi(\boldsymbol{\theta})$ the pdf of a distribution of interest, usually referred to as the *target* distribution. It is often impractical to sample $\pi(\boldsymbol{\theta})$ directly, so we are going to draw samples from a *proposal* distribution and weight them appropriately according to the principle of IS [1]. If $q(\boldsymbol{\theta})$ denotes the proposal pdf, then the idea is to draw samples,

$$\boldsymbol{\theta}^{(i)} \sim q(\boldsymbol{\theta}),\ i = 1, \cdots, M \tag{2}$$

and assign each one a normalized importance weight, $w^{(i)}$, computed as

$$w^{(i)*} \propto \frac{\pi(\boldsymbol{\theta}^{(i)})}{q(\boldsymbol{\theta}^{(i)})}, \quad w^{(i)} = \frac{w^{(i)*}}{\sum_{i=1}^M w^{(i)*}}, \quad i = 1, \cdots, M \ . \tag{3}$$

Equations (2) and (3) together constitute the standard IS algorithm.

From the sample set $\Theta^M = \{\boldsymbol{\theta}^{(i)}\}_{i=1}^M$ and their associated weights, one can build a discrete random measure,

$$\pi^M(d\boldsymbol{\theta}) = \sum_{i=1}^M w^{(i)}\delta_{\boldsymbol{\theta}^{(i)}}(d\boldsymbol{\theta}), \tag{4}$$

where $\delta_{\boldsymbol{\theta}^{(i)}}$ is the unit delta measure located at $\boldsymbol{\theta} = \boldsymbol{\theta}^{(i)}$, that allows to approximate the expectation of any integrable function f with respect to $\pi(\boldsymbol{\theta})$ as

$$\mathbb{E}_{\pi(\boldsymbol{\theta})}\left[f(\boldsymbol{\theta})\right] \approx \sum_{i=1}^M w^{(i)} f(\boldsymbol{\theta}^{(i)}). \tag{5}$$

The efficiency of any IS algorithm (roughly given by the number of samples that are required to attain a certain level of performance) depends to a great extent on the choice of the proposal function, $q(\boldsymbol{\theta})$. However, the asymptotic convergence of the above approximation when $M \to \infty$ is guaranteed as long as $0 < \frac{\pi(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} < \infty$ for every $\boldsymbol{\theta}$ [1].

Weight degeneracy and transformation of the weights: The degeneracy of the IWs refers to the situation in which only a few samples hold significant IWs, while the vast majority have weights close to zero. This clearly reduces the efficiency of the method since samples with negligible IWs barely contribute to the approximations in (4) or (5).

Following [4], we propose to alleviate this problem by applying a non-linear transformation over the unnormalized IWs, $w^{(i)*}$, aimed at decreasing their variance. To be specific, the (unnormalized) transformed IWs (TIWs) are obtained as $\bar{w}^{(i)*} = \varphi_{\Theta^M}\left(w^{(i)*}\right), i = 1, \cdots, M$, where $\varphi_{\Theta^M}(\cdot)$ is a real positive function that depends on the whole sample set, Θ^M, and their weights.

Different choices for the non-linear transformation, $\varphi_{\Theta^M}(\cdot)$, are possible. The asymptotic convergence results in [4] are specifically referred to the "clipping" transformation, which is also the one we consider here. It consists in setting the $M_T < M$ highest IWs to a common value. More formally, let us consider a permutation, $i_1, \cdots, i_M$ of the indices in $\{1, \cdots, M\}$ such that $w^{(i_1)*} \geq \cdots \geq w^{(i_{M_T})*} \geq \cdots \geq w^{(i_M)*}$. Then, the unnormalized TIWs,

$\bar{w}^{(i)*}$, are computed from the original IWs, $w^{(i)*}$, as

$$\bar{w}^{(i)*} = \varphi_{\Theta^M}\left(w^{(i)*}\right) = \min\left(w^{(i)*}, w^{(i_{M_T})*}\right), \tag{6}$$

and their normalized counterparts are

$$\bar{w}^{(i)} = \frac{\bar{w}^{(i)*}}{\sum_{i=1}^{M} \bar{w}^{(i)*}}, i = 1, \cdots, M. \tag{7}$$

The IS algorithm with TIWs is described by equations (2), (3), (6) and (7). We refer to it as nonlinear IS (NIS).

Simulations: We numerically asses the convergence of the NIS algorithm when the sample size is finite and, most importantly, when using a non-iterative scheme (unlike in [4]). We apply the algorithm to estimate the location of the modes of a Gaussian mixture model (GMM) given a set of conditionally independent and identically distributed (i.i.d.) observations. In particular, we consider the GMM

$$p(y|\boldsymbol{\theta}) = \rho_1 \mathcal{N}\left(y|\theta_1, \sigma^2\right) + \rho_2 \mathcal{N}\left(y|\theta_2, \sigma^2\right) + (1 - \rho_1 - \rho_2)\mathcal{N}\left(y|\theta_3, \sigma^2\right), \tag{8}$$

where $\boldsymbol{\theta} = [\theta_1, \theta_2, \theta_3]^\top = [0, 2, 4]^\top$ is a vector encompassing the unknown means of the mixture components. The remaining parameters of the model are assumed to be known, and set to $\rho_1 = 0.2$, $\rho_2 = 0.3$ and $\sigma^2 = 1$.

We assume independent Gaussian prior distributions for the unknown means, namely

$$p(\boldsymbol{\theta}) = \mathcal{N}\left(\theta_1|1, 10\sigma^2\right)\mathcal{N}\left(\theta_2|1, 10\sigma^2\right)\mathcal{N}\left(\theta_2|1, 10\sigma^2\right). \tag{9}$$

Given a collection of N observations, $\mathbf{y} = \{y_1, y_2, \cdots, y_N\}$, drawn from the GMM in (8), we aim at approximating the posterior $\pi(\boldsymbol{\theta}|\mathbf{y})$. In order to do so, we draw samples from a proposal function, which is selected to match the prior, i.e.,

$$\boldsymbol{\theta}^{(i)} \sim q(\boldsymbol{\theta}) = p(\boldsymbol{\theta}), i = 1, \cdots, M. \tag{10}$$

Then, at the sight of equations (3) and (1), the unnormalized IWs are computed as

$$w^{(i)*} \propto \frac{\pi(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} \propto p(\mathbf{y}|\boldsymbol{\theta}) = \prod_{i=1}^{N} p(y_i|\boldsymbol{\theta}), \tag{11}$$

where we have used the fact that the observations in $\mathbf{y}$ are conditionally independent given the state. TIWs can then be computed as indicated in (6) and (7).

In this setting we compare the performance of the standard IS and the novel NIS schemes. Every result in this section is averaged over a certain number of independent simulation runs, R, of the appropriate algorithm for a fixed set of observations $\mathbf{y}$. We refer to each of these runs as an MC realization, and we set $R = 1,000$. Additionally, L independent realizations of the observations, $\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_L$, are considered. Overall, each algorithm is run LR times.

Figure 1 shows the performance of the standard IS and NIS algorithms in terms of the bias of the posterior-mean estimators, $\hat{\boldsymbol{\theta}}_{NIS} = \sum_{i=1}^{M} \bar{w}^{(i)}\boldsymbol{\theta}^{(i)}$ and $\hat{\boldsymbol{\theta}}_{IS} = \sum_{i=1}^{M} w^{(i)}\boldsymbol{\theta}^{(i)}$, as the number of samples, M, grows. For every value of M, the number of samples whose weights are clipped is $M_T = \log(M)$. On the other hand, the number of observations is in every case $N = 1,000$. The bias is computed as

$$\text{Bias} = \frac{1}{L}\sum_{l=1}^{L}\left\|\frac{1}{R}\sum_{r=1}^{R}\hat{\boldsymbol{\theta}}_{l,r} - \boldsymbol{\theta}\right\|,$$

where $\hat{\boldsymbol{\theta}}_{l,r}$ is the posterior-mean estimate of $\boldsymbol{\theta}$ computed for the l-th realization of the observations, $\mathbf{y}_l$, during the r-th MC realization for either the NIS or IS algorithms. Additionally, the figure also provides information about the degeneracy of the weights. The color of every marker indicates, according to the color bar on the right, the maximum weight among all the samples. It can be seen that the bias attained by the standard IS estimator is always below

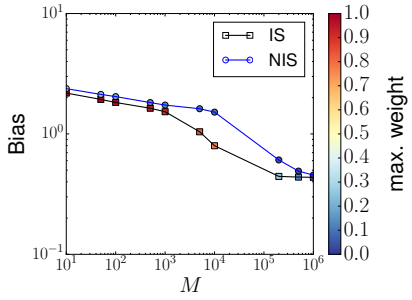

Fig. 1. Bias averaged over $L = 75$ observations realizations.

the bias of the NIS estimator. This is because the transformation of the weights applied to obtain the TIWs introduces a distortion in the approximation of the probability measure $\pi(\boldsymbol{\theta})d\boldsymbol{\theta}$. However, it is clear from the figure that both algorithms converge to the same estimate when the number of samples is large as predicted by the asymptotic analysis in [4]. Also notice that in the standard IS algorithm, the maximum weight is close to 1 whenever the number of samples is below $M = 10,000$. Hence, most of the probability mass is concentrated in a single sample, and the remaining samples are, for all practical purposes, irrelevant. This is specifically avoided in the NIS algorithm, where the maximum weight is below 0.2 for most values of M. Avoiding a single sample garnering most of the probability mass is important in connection to the variance of the resulting estimators. This is illustrated in Figure 2 that shows the estimator variance, computed as $\text{Variance} = \frac{1}{L}\sum_{l=1}^{L}\text{trace}\left\{\frac{1}{R}\sum_{r=1}^{R}\left(\hat{\boldsymbol{\theta}}_{l,r} - \bar{\hat{\boldsymbol{\theta}}}_l\right)\left(\hat{\boldsymbol{\theta}}_{l,r} - \bar{\hat{\boldsymbol{\theta}}}_l\right)^\top\right\}$ with $\bar{\hat{\boldsymbol{\theta}}}_l = \frac{1}{R}\sum_{r'=1}^{R}\hat{\boldsymbol{\theta}}_{l,r'}$. The figure shows that, for low-to-medium sample sizes, the variance of the estimates computed using the NIS method is considerably lower than the variance of the standard IS estimates.

A common metric for comparing the performance of different algorithms is the mean square error (MSE). Our simulations indicate that, e.g., for $M = 1,000$ the (average) MSE attained by the standard IS algorithm is 6.21 while that achieved by the NIS scheme is 3.82.

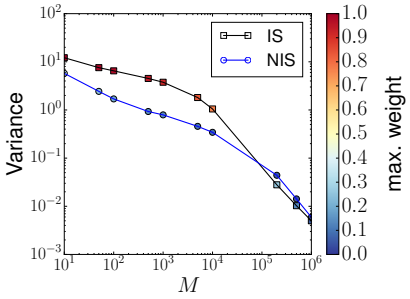

Fig. 2. Variance averaged over $L = 75$ observations realizations.

In the last experiment, we explore the impact on the performance of the NIS scheme of both the number of observations, N, and the number of clipped weights, M_T. Notice that when $M_T = 1$, the resulting algorithm is the standard importance sampler. Figure 3 shows that, for any number of observations N, the bias (left) increases along M_T, whereas the variance (right) decreases. However, there is an elbow in the curves for the variance, which suggests that increasing the value of M_T above 10 does not yield any benefits in terms of variance while the bias keeps increasing linearly. Another remarkable result is that, as N, grows, the bias decreases and the variance increases. This is due to the target pdf concentrating in an ever smaller region as the number of observations, N, grows, which, in turn, makes sampling more difficult. In such a case, the benefits stemming from using NIS are more obvious. This can be seen by comparing, e.g., the variance when $M_T = 1$ (plain IS) and $M_T = 5$ for $N = 1,000$ observations.

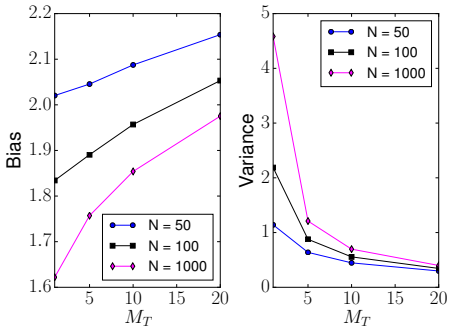

Fig. 3. Bias (Left) and variance (Right) obtained by the NIS algorithm for $M = 500$. The results are averaged over $L = 200$ observations realizations.

Conclusion: We have investigated the benefits of applying a nonlinear transformation to the IS weights in order to alleviate the well-known degeneracy problem. Our computer simulations show that, while both the IS and NIS schemes converge to the same approximations when the number of samples is large enough, the estimators computed via the NIS method attain an advantageous variance/bias trade-off that often results in a better practical performance.

Acknowledgment: This work was supported by *Ministerio de Economía y Competitividad* of Spain (projects TEC2012-38883-C02-01 and TEC2015-69868-C2-1-R) and the Office of Naval Research Global (award no. N62909- 15-1-2011).

M. A. Vázquez and J. Míguez (*Departamento de Teoría de la Señal y Comunicaciones, Universidad Carlos III de Madrid, Spain*)

E-mail: {mvazquez,jmiguez}@tsc.uc3m.es

References

1 C. P. Robert, G. Casella, Monte Carlo Statistical Methods, Springer, 2004.
2 T. Bengtsson, P. Bickel, B. Li, et al., Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems, in: Probability and statistics: Essays in honor of David A. Freedman, Institute of Mathematical Statistics, 2008, pp. 316–334.
3 O. Cappe, A. Guillin, J. Marin, C. Robert, Population Monte Carlo, Journal of Computational and Graphical Statistics 13 (4) (2004) 907–929.
4 E. Koblents, J. Míguez, A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models, Statistics and Computing 25 (2) (2015) 407–425.